



IN THE UNITED STATES PATENT & TRADEMARK OFFICE

In re Application of:  
**Timothy Jerry Schimke, et al.**

Serial No.: **09/819,122**

Filed: **27 March 2001**

Title: **METHOD AND SYSTEM FOR  
ACCURATELY DETERMINING A  
DEVICE LOCATION IN AN  
ARBITRATED LOOP**

§  
§  
§  
§  
§  
§  
§  
§  
§  
§

Attorney Docket No.: **RC920000254US1**

Examiner: **CASIANO, ANGEL L.**

Group Art Unit: **2182**

**DECLARATION OF PRIOR INVENTION IN THE UNITED STATES  
TO OVERCOME CITED PATENT UNDER 37 C.F.R. § 1.131**

Mail Stop Non-Fee Amendment  
Commissioner for Patents  
P. O. Box 1450  
Alexandria, Virginia 22313-1450

Sir:

**PURPOSE OF DECLARATION**

This Declaration is to establish completion of the invention in this Application in the United States at a date prior to 01 February 2000, which is the effective filing date of United States Patent No. 6,678,839 B2, issued to Mori on 13 January 2004, and cited by the Examiner.

The persons making this Declaration are the only co-inventors of the invention related to the above captioned patent application.

### **FACTS AND DOCUMENTARY EVIDENCE**

To establish the date of completion of this Application, Applicants submit herewith International Business Machines Corporation Invention Disclosure Form ROC8-2000-0359, titled "Method and Apparatus for Accurately Determining Device Location in the Presence of Hardware Faults," last updated on 02 June 2000, from which it can be seen at the bottom of Page 5 of the disclosure that the invention in this Application was completed at least by 26 July 1999, a date which is earlier than the effective date of the reference.

Applicant asserts that the content of the attached Invention Disclosure is sufficiently complete to form a constructive reduction to practice.

### **DILIGENCE**

In the event the Examiner determines that the Invention Disclosure Form submitted herewith lacks sufficient completeness to constitute a constructive reduction to practice, Applicants hereby assert diligence in the completion of this invention and the filing thereof from a time prior to 26 July 1999 continuously up to the date of the constructive reduction to practice, which occurred on the filing of the present application on 27 March 2001. At all times between 26 July 1999 and 27 March 2001, Applicants were diligent in that the Invention Disclosure was submitted to the Intellectual Property Law Department of International Business Machines Corporation and was processed as required by corporate policies. That Disclosure was then evaluated by a Patent Committee and filing of the parent application was authorized. Thereafter, the Invention Disclosure was submitted to counsel who was diligent in preparation of the present patent application which was filed on 27 March 2001.

### **TIME OF PRESENTATION OF THE DECLARATION**

This Declaration is submitted prior to a Final Rejection.

## DECLARATION

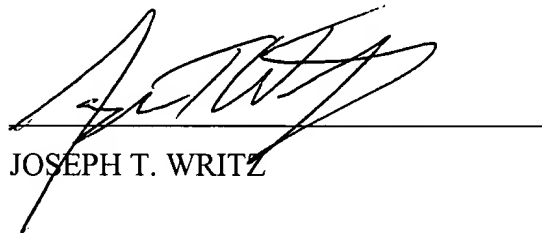
As a person signing below, I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful, false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the U.S.C., and that such willful, false statements may jeopardize the validity of the Application or any patent issued thereon.



TIMOTHY JERRY SCHIMKE



RICK WECKWERTH



JOSEPH T. WRITZ

Roc9-2000-0254

9/11/2000



## Disclosure ROC8-2000-0359

**Created By:** Tim Schimke **Created On:** 05/05/2000 04:44:13 PM

**Last Modified By:** Tim Schimke **Last Modified On:** 06/02/2000 08:31:09 AM

\*\*\* IBM Confidential \*\*\*

Required fields are marked with the asterisk (\*) and must be filled in to complete the form .

### Summary

Status	Under Evaluation
Processing Location	ROC
Functional Area	2H3 - SD - Jon Claeys (DASD Attach) (Hefner)
Attorney/Patent Professional	James R Nock/Rochester/IBM
IDT Team	Steven DeFoster/Rochester/IBM; Shawn Lambeth/Rochester/IBM; Chuck Graham/Rochester/IBM; Paul Movall/Rochester/IBM; Robert Williams/Rochester/IBM; Fred Huss/Rochester/IBM; Gary Mulford/Rochester/IBM; Andy Kulich/Rochester/IBM; Jim Tilbury/Rochester/IBM; Gene Van Grinsven/Rochester/IBM; AL Yanes/Rochester/IBM; Brian Bakke/Rochester/IBM; Jerry Grabowski/Rochester/IBM; Gary Reuland/Rochester/IBM; Don Ziebarth/Rochester/IBM; James R Nock/Rochester/IBM
Submitted Date	06/01/2000 04:37:08 PM
Owning Division	SD
Select	
	To calculate a PVT score, use the 'Calculate PVT' button.
Incentive Program	
Lab	
Technology Code	

### Inventors with Lotus Notes IDs

Inventors: Tim Schimke/Rochester/IBM, Rick Weckwerth/Rochester/IBM, Joe Writz/Rochester/IBM

Inventor Name > denotes primary contact	Inventor Serial	Div/Dept	Manager Serial	Manager Name
> Schimke, Timothy J. (Tim)	294389	7T/HTTA	654681	Krsnak, Roger D
Weckwerth, Rick	532309	7T/CC7A	045892	Rink, L. C. (Luz)
Writz, Joseph T (Joe)	525357	7T/CC7A	045892	Rink, L. C. (Luz)

### Inventors without Lotus Notes IDs

### IDT Selection

<b>IDT Team:</b> Steven DeFoster/Rochester/IBM Shawn Lambeth/Rochester/IBM Chuck Graham/Rochester/IBM Paul Moval/Rochester/IBM Robert Williams/Rochester/IBM Fred Huss/Rochester/IBM Gary Mulford/Rochester/IBM Andy Kulich/Rochester/IBM Jim Tilbury/Rochester/IBM Gene Van Grinsven/Rochester/IBM AL Yanes/Rochester/IBM Brian Bakke/Rochester/IBM Jerry Grabowski/Rochester/IBM Gary Reuland/Rochester/IBM Don Ziebarth/Rochester/IBM James R Nock/Rochester/IBM	<b>Attorney/Patent Professional:</b> James R Nock/Rochester/IBM
---	--

**Response Due to IP&L : 07/02/2000**

## Main Idea

### \*Title of disclosure (in English)

Method and Apparatus for Accurately Determining Device Location in the Presence of Hardware Faults

### \*Idea of disclosure

1. Describe your invention, stating the problem solved (if appropriate), and indicating the advantages of using the invention.

It is necessary to accurately determine the exact location of peripheral devices for serviceability reasons. The failing device must be identifiable so that it may be replaced upon failure, which implies a need to specify the location down to a single slot within a system unit or expansion tower. The location of the device is ordinarily determined by mapping the logical address the device communicates with on the loop to the physical location of the device. The device chooses a logical address based upon a value presented to it by the backplane connector (each peripheral device location has a unique value). A problem occurs if there is a fault in the peripheral device or backplane such that the logical address the device uses does not map correctly to the actual physical location.

Our invention is an mechanism for detecting an incorrect mapping caused by a failure in the backplane or peripheral device(s) when FC-AL devices in a private loop are used in conjunction with Port Bypass Circuits (PBCs), and also determining what the actual correct mapping is in the presence of the failure(s). Our invention also helps to identify the faulting piece of hardware so that it may be replaced.

2. How does the invention solve the problem or achieve an advantage,(a description of "the invention", including figures inline as appropriate)?

Our invention uses the Port Bypass Circuits (PBCs) in conjunction with information obtained when talking to peripheral devices on the loop to detect cases where the mapping from logical address to physical address is incorrect. The correct mapping is determined, and faulty pieces of hardware are identified for replacement. The description of the invention will include the following steps, each of which is further elaborated upon below:

- A) Brief description of FC-AL Technology and the Need for PBCs.
- B) Description of FC-AL Logical Addressing Method.
- C) Description of Enhanced Logical Address(ALPA) to Physical Address Mapping.

A) Description of FC-AL Technology and Need for PBCs.

FC-AL is a node-to-node loop technology where the failure of a single node on the loop causes the entire loop to fail. Empty slots within a backplane are handled with a simple bypass circuit which

automatically bypass a device slot (leaving the loop functioning) if the device is absent or requests to be bypassed. Devices can take themselves off of the loop using the bypass signal if they are not ready to communicate, believe they are faulty, or receive an FC-AL in-band loop primitive instructing them to exit the loop. These methods rely on the device being well-behaved even in the presence of failures to prevent loop failure. PBCs are an additional mechanism which may be used to prevent loop failure. PBCs are an enhanced bypass circuit which have an additional input that allows the signals from the device to be overridden, enabling a misbehaving device to be isolated. Each device slot has its own PBC. All communication with the PBCs is done out-of-band (i.e. not using the FC-AL loop itself, instead using an alternative mechanism such as an I<sup>2</sup>C bus that works even if the loop itself is non-functional). If the loop is non-functional, the PBCs can be used to determine the identity of a misbehaving device and bypass it allowing the loop to recover.

#### B) Description of FC-AL Logical Addressing Method.

Each node on an FC-AL loop is assigned an 8 bit ALPA. The ALPA is the logical address used to communicate with the node, and ALPAs are unique within a given loop. Peripheral devices receive a SEL\_ID through the backplane connector. The SEL\_ID is a 7 bit value that may be algorithmically converted to an 8 bit "preferred ALPA". When loop initialization occurs after initial power-up, each node attempts to obtain its "preferred ALPA". If there are no duplicate ALPA values among the nodes on the loop, then each node's actual ALPA will be equal to its "preferred ALPA". If collisions do exist, then there is a defined mechanism through which all nodes are assigned unique ALPA values. In this case the actual ALPA will be different from the "preferred ALPA" for one or more nodes.

NOTE: The full ALPA assignment sequence includes two additional steps which precede the above steps. The first is "fabric assigned", and this is not applicable in this case because it is a private loop. The second is "previously-assigned ALPA", and would be the ALPA assigned the last time loop initialization occurred during this power-up sequence for the node. The "previously-assigned ALPA" is applicable, and this has implications in some cases.

The normal case is for each node on the loop to have a unique SEL\_ID value, which will lead to each node receiving its "preferred ALPA". This ALPA, in conjunction with knowledge of the physical packaging, may then be used by the initiator or host driver software to determine the physical location of each device by using the mapping from SEL\_ID to ALPA. If there is a fault (bad pin, broken wire, etc.) in either the node or the backplane, the SEL\_ID used by the node may not be what was intended. This may result in the ALPA actually used on the loop for the faulty node to no longer be the "preferred ALPA". Other non-faulty nodes on the loop may also not receive their "preferred ALPA" because multiple nodes may now be attempting to obtain the same ALPA. In general, a single fault can cause none, some, or all nodes not to receive their "preferred ALPA" depending on the configuration, the SEL\_ID values, and the error encountered.

Because the mapping from ALPA to physical location identifies the correct location only if the nodes receive their "preferred ALPA", the system will use incorrect data if some faults are present in the system. Note that the peripheral devices may still be functioning normally even in the presence of faults, so that there is no indication of problems until sometime later when maintenance is attempted. A single fault may result in all nodes being tagged with incorrect physical locations. This will complicate maintenance of a failing resource, and may require a more disruptive recovery action to occur. Systems built upon SCSI devices have a similar concern. However the SCSI case is typically less severe because in the SCSI case there is no dynamic reassignment of logical addresses so failures typically only affect the faulting device or the faulting device plus the device whose address he now conflicts with.

#### C) Description of Enhanced Logical Address(ALPA) to Physical Address Mapping.

Our invention is a method to determine more accurately the physical location of each logical address. Ordinarily the PBCs are used solely to bypass a device that is preventing proper loop operation. We use the PBCs as a tool to determine the exact mapping between physical addresses and logical

addresses. The sequence of steps is used at initial power-on time, or at other times as needed such as after a concurrent maintenance action.

The following steps are used assuming there are no failing devices with PBCs already enabled (this is not a requirement, but does simplify the explanation):

1. The out-of-band communication is done with the PBCs to determine the set of valid physical addresses for peripheral devices, i.e. the set of values that will be provided to the devices via the backplane connector. The initiator will use a preferred ALPA that does not map to any of these physical addresses.
2. All PBCs are enabled such that only the initiator is present on the loop. The loop is then brought up and the set of ALPAs examined. Only the initiator's ALPA is present since all other devices have been bypassed. This determines the ALPA of the initiator, and we can validate that the initiator's ALPA does not map to any of the valid physical addresses.
3. The next step is to determine the logical address used by the peripheral device inserted into each physical slot. This is done by enabling all PBCs except for the initiator and one other slot, and determining the ALPAs used and the WWID of that peripheral device. The dynamic nature of an ALPA generated from a SEL\_ID is not a concern at this time because the number of entries on the loop has been constrained to only the initiator and a single peripheral device so that there will not be any ALPA collisions. This step is repeated for each physical slot. After the completion of all physical slots a table indexed by physical location is constructed containing the ALPA and WWID for each physical slot. This table is denoted as table A.
4. The next step is to validate each entry of table A. The fault isolation procedure is different at initial power-on time than at later times because later loop initialization may have nodes using the "last-known ALPA" value.
  1. Initial power-on time: The actual ALPA in this case for each entry in table A was determined using what the device believed its SEL\_ID value was because the first time the device entered the loop after initial power-on it was the only device besides the initiator. Any entry in which the ALPA generated from the physical address (i.e. SEL\_ID) is not the same as the actual ALPA indicates a fault in either the backplane or the that peripheral device. These faults are reported to the host system so that maintenance can be performed.
  2. Later times: A history of the last execution of this sequence of steps is kept to properly identify the faulting device(s). An entry in table A for a node which existed last time (i.e. same WWID) is skipped because it used the "previously-assigned ALPA" value and any needed fault reporting would have occurred on early execution of this sequence. Each entry which did not exist last time (i.e. no matching WWID in the last table A) is checked. These devices most likely entered the loop dynamically (i.e. concurrently added) when other devices were already on the loop. In this case the device attempted to use its SEL\_ID value to obtain an ALPA, but would not have been successful if another node already had it. Because of this, a fault is reported only if the ALPA generated from the physical address (i.e. SEL\_ID) is different than the actual ALPA, and there is no other entry in table A which does have the ALPA mapped from the physical address for this node. (NOTE: It is possible (but unlikely) for a fault to go undetected if a device is added in the presence of other faults previously reported. This fault would have to be present in the device slot or device being added, and would be found the next time the system goes through an initial power-up sequence.)
5. The final step is to remove the PBC bypass on all physical slots. Optionally the device slots with faults present can be bypassed until maintenance is completed, this may simplify error recovery procedures in the host system. A second table consisting of ALPAs and WWIDs for each ALPA is constructed based upon the loop configuration with all devices present. The physical location is not yet filled in, that will be done in step 6. The second table is denoted as table B.

6. The physical location for each entry in table B is determined using table A. For each entry in table B, the physical location is obtained from the entry in table A containing the matching WWID. The ALPA is not used because a given device's ALPA may have been remapped to a different value if there were multiple devices attempting to obtain the same ALPA. This might occur if multiple devices thought they had the same SEL\_ID because of a hardware fault.

**NOTE:** If all physical slots with identified faults were bypassed in step 5, then the physical location retrieved out of table A will be the same as the physical location generated algorithmically from ALPA.

**Additional Notes:**

1. If there are PBC(s) already enabled for failing device(s), the exact same steps may be used except that any physical slot being bypassed because of a device fault remains bypassed for all steps.
2. The initial power-on sequence of steps can be used at later times if there is a method to force the nodes not to use the "previously-assigned ALPA" value when obtaining the ALPA during loop initialization. Examples of such method might be forcing the nodes back through power-on, a vendor-unique method to make the node "forget" he already had obtained an ALPA previously, a mode-pin disabling all use of the "previously-assigned ALPA", or some other mechanism.

3. If the same advantage or problem has been identified by others (inside/outside IBM), how have those others solved it and does your solution differ and why is it better?

This problem has long existed when using SCSI peripheral devices and is also seen when using FC-AL peripheral devices. Ordinarily only the logical address the device responds to is used to calculate the expected physical location. The cases where this mapping generates the incorrect physical location due to a fault in the device or backplane are most often ignored. Another method which may be used is for the device to communicate the ALPA desired and the ALPA actually received to the initiator via a SCSI inquiry page or another mechanism. This is an improvement but not a complete solution because it is reliant on the device correctly determining the ALPA he was supposed to obtain, and there may have been a fault such that the device was attempting to obtain an incorrect ALPA.

Because other schemes do not accurately identify the physical location in all cases, it results in a more drastic recovery action being required for those cases where the mapping was incorrect. This recovery action may not be possible concurrent with normal system operation resulting in significant customer impact. Possible recovery actions may include a save/restore of the system or lengthy reconfiguration of the peripheral devices.

Our invention detects cases where the mapping between logical address and expected physical location is incorrect prior to device failure, and also determines the correct mapping in the presence of the device or backplane fault. This allows normal concurrent maintenance procedures to be used to repair the failing device and minimizes customer impact.

4. If the invention is implemented in a product or prototype, include technical details, purpose, disclosure details to others and the date of that implementation.

High level design, architecture, and specification reviews were completed. The project requiring the function provided by our invention was redefined prior to actual implementation of our invention.

**\*Critical Questions ( Questions 1 - 7 must be answered)**

**\*Question 1**

On what date was the invention workable? 07/26/99 Please format the date as MM/DD/YYYY (Workable means i.e. when you know that your design will solve the problem)



<b>*Question 2</b>		<input type="radio"/> Yes
Is there any planned or actual publication or disclosure of your invention to anyone outside IBM?		<input checked="" type="radio"/> No
If yes, Enter the name of each publication or patent and the date published below.		
Publication/Patent:		
Date Published or Issued:		
Are you aware of any publications, products or patents that relate to this invention?		<input type="radio"/> Yes
		<input checked="" type="radio"/> No
If yes, Enter the name of each publication or patent and the date published below.		
Publication/Patent:		
Date Published or Issued:		

<b>*Question 3</b>		<input type="radio"/> Yes
Has the subject matter of the invention or a product incorporating the invention been sold, used internally in manufacturing, announced for sale, or included in a proposal?		<input checked="" type="radio"/> No
Is a sale, use in manufacturing, product announcement, or proposal planned?		<input type="radio"/> Yes
		<input checked="" type="radio"/> No
If Yes, identify the product if known and indicate the date or planned date of sale, announcements, or proposal and to whom the sale, announcement or proposal has been or will be made.		
Product:		
Version/Release:		
Code Name:		
Date:		
To Whom:		
If more than one, use cut and paste and append as necessary in the field provided.		

<b>*Question 4</b>		<input type="radio"/> Yes
Was the subject matter of your invention or a product incorporating your invention used in public, e.g., outside IBM or in the presence of non-IBMers?		<input checked="" type="radio"/> No
If yes, give a date. <b>Please format the date as MM/DD/YYYY</b>		

<b>*Question 5</b>		<input type="radio"/> Yes
Have you ever discussed your invention with others not employed at IBM?		<input checked="" type="radio"/> No
If yes, identify individuals and date discussed. Fill in the text area with the following information, the names of the individuals, the employer, date discussed, under CDA, and CDA #.		

<b>*Question 6</b>		<input type="radio"/> Yes
Was the invention, in any way, started or developed under a government contract or project?		<input checked="" type="radio"/> No
		<input type="radio"/> Not sure
If Yes, enter the contract number		

<b>*Question 7</b>		<input type="radio"/> Yes
Was the invention made in the course of any alliance, joint development or other contract activities?		<input checked="" type="radio"/> No
		<input type="radio"/> Not Sure
If Yes, enter the following :Name of Alliance, Contractor or Joint Developer		
Contract ID number		
Relationship contact name		
Relationship contact E-mail		
Relationship contact phone		

**Question 8**

Have you submitted, or are you aware of, any related disclosure submission?

☐ Yes  
☒ No

If Yes, please provide the title and docket or disclosure number below:

**Question 9**

What type of companies do you expect to compete with inventions of this type? *Check all that apply.*

- ☒ Manufacturers of enterprise servers
- ☒ Manufacturers of entry servers
- ☒ Manufacturers of workstations
- ☒ Manufacturers of PC's
- ☐ Non-computer manufacturers
- ☐ Developers of operating systems
- ☐ Developers of networking software
- ☐ Developers of application software
- ☒ Integrated solution providers
- ☐ Service providers
- ☐ Other (Please specify below)

**Patent Value Tool (Optional - this may be used by the inventor and attorney to assist with the evaluation)**

(The Patent Value tool can be used by you or the evaluation team to determine the potential licensing value of your invention.)

The **Patent Value Tool** has not yet been used to calculate a score.

**Post Disclosure Text & Drawings**

Enter any additional information relating to this disclosure below:

---

(Form Revised 12/17/97)